

Stochastic Feature Mapping for PAC-Bayes Classification

Xiong Li, and Tai Sing Lee, and Yuncai Liu, *Senior Member, IEEE*,

Abstract—Probabilistic generative modeling of data distributions can potentially exploit hidden information which is useful for discriminative classification. This observation has motivated the development of approaches that couple generative and discriminative models for classification. In this paper, we propose a new approach to couple generative and discriminative models in an unified framework based on PAC-Bayes risk theory. We first derive the *model-parameter-independent stochastic feature mapping* from a practical MAP classifier operating on generative models. Then we construct a linear stochastic classifier equipped with the feature mapping, and derive the *explicit PAC-Bayes risk bounds* for such classifier for both supervised and semi-supervised learning. Minimizing the risk bound, using an EM-like iterative procedure, results in a new posterior over hidden variables (E-step) and the update rules of model parameters (M-step). The derivation of the posterior is always feasible due to the way of equipping feature mapping and the explicit form of bounding risk. The derived posterior allows the tuning of generative models and subsequently the feature mappings for better classification. The derived update rules of the model parameters are same to those of the uncoupled models as the feature mapping is model-parameter-independent. Our experiments show that the coupling between data modeling generative model and the discriminative classifier via a stochastic feature mapping in this framework leads to a general classification tool with state-of-the-art performance.

Index Terms—stochastic feature mapping; PAC-Bayes risk bound; hybrid generative-discriminative classification

I. INTRODUCTION

Discriminative models designed to find decision boundaries among different classes are state-of-the-art tools for classification, while probabilistic generative models seeking to model data distributions are adept in exploiting hidden information, in dealing with structured data (e.g. protein sequence with variable length) and in solving nonlinear classification problems using maximum a posterior (MAP) classifier. The complementarities of the two paradigms have been investigated [19], [8], resulting in several promising works [3], [21], [23], [9]. The above observations have emerged from these works in the context of classification: (1) generative models provide feature mappings that simultaneously exploit hidden information, and transform structured data into a fixed dimensional feature; (2) discriminative models find an optimum decision boundaries in such a feature space under specified criterion.

Generative score space methods [3], [9], [14] are motivated by the above observations. These methods derive feature mappings from the log likelihood (or its lower bound) of generative

models. These feature mappings are measures over models $P(\mathbf{x}, \mathbf{h}, \theta)$, taking the form of $E_{P(\mathbf{h}|\mathbf{x})}[\phi(\mathbf{x}, \mathbf{h}, \theta)]$ where ϕ is a function over the observed variable \mathbf{x} and the hidden variable set \mathbf{h} . They map observed and hidden variables into a vector of score, which are then used as features by classifiers. These methods exploit the superior abilities of generative models in exploiting hidden information and dealing with structured data. However, in these methods, generative models are isolated from the classification process and there is no principled way to tune the generative models as well as the feature mapping to improve classification. It is desirable to develop a mechanism that can couple the classifier to the generative models to allow fine-tuning of the feature mapping.

Maximum entropy discrimination [19] provides yet another framework to exploit generative models for classification under the large margin principle. This framework, however, requires deliberately choosing conjugate priors for parameters of the generative models, which limits its application to complex models. In addition, the VC risk bound [11] utilized by this method is generally loose in comparison with the PAC-Bayes bounds [13], [5], [2]. Also, there are some other efforts [21], [23] made to couple generative and discriminative models for classification. However, these methods provide no explicit feature mapping which is useful in real applications. Further, they requires re-formulating the update rules of the parameters of generative models, which is typically complex.

This paper proposes an approach based on the PAC-Bayes theory [13], [5], [2] to integrate the complementary strengths of generative and discriminative models. Using the linear form of a practical MAP classifier operating on generative models, we derive the *model-parameter-independent stochastic feature mapping*. By the feature mapping, we meant that the feature used for classifications is a function of the input data and the hidden variables of the generative models. This is distinct from the current methods [3], [9], [14] which map a data point to a feature deterministically. Then we construct a *stochastic classifier*, a Gibbs classifier operating on the derived feature mapping, and derive explicit PAC-Bayes risk bounds for such a classifier. By minimizing the risk bound using an EM-like iterative procedure, we derive the *posterior over the hidden variables* (E-step) and the *update rules of model parameters* (M-step). The derivation is always feasible due to the way of equipping feature mapping and the form of bounding risk. The posterior provides a bridge that allows the classifier to tune the generative models and subsequently the feature mapping for classification. The update rules of model parameters are quite simple – essentially same to those of the uncoupled models as the feature mapping is model-parameter-independent.

Xiong Li and Yuncai Liu are with the Department of Automation, Shanghai Jiao Tong University, Shanghai, China; Tai Sing Lee and Xiong Li are with the Computer Science Department, Carnegie Mellon University, PA, USA.
E-mail: {lixiong, whomliu}@sjtu.edu.cn, tai@cs.cmu.edu

II. FROM MAP CLASSIFIER TO STOCHASTIC FEATURE MAPPING

In this section, for exponential family generative models, we derive the linear form (Eq. (4)) of the MAP classifier (Eq. (3)) based on the variational approximation (Eq. (2)), and use that to derive a stochastic feature mapping. The derived feature mapping is functioning similar to [3], [9], [14]. Consider the binary classification problem that assigns labels $y \in \{-1, +1\}$ to examples $\mathbf{x} \in \mathbb{R}^d$. Let $P(\mathbf{x}|\theta_y)$ be the class-conditional distributions over \mathbf{x} ; $P(y)$ be the prior of labels. The decision rule of the MAP classifier is $\hat{y} = \max_y P(\mathbf{x}|\theta_y)P(y)$, which is equivalent to $\hat{y} = \text{sign}(L(\mathbf{x}; \Theta))$ where $\text{sign}(a) = +1$ if $a > 0$ and $\text{sign}(a) = -1$ otherwise, and the discriminant function:

$$L(\mathbf{x}, \Theta) = \log P(\mathbf{x}|\theta_+) - \log P(\mathbf{x}|\theta_-) + b \quad (1)$$

where subscripts $+$, $-$ are the shorts of $+1$, -1 ; $\Theta = \{\theta_-, \theta_+, b\}$; $b = \log P(y = +1) - \log P(y = -1)$. When $P(\mathbf{x}|\theta_y)$ is modeled by a generative model $P(\mathbf{x}, \mathbf{h}|\theta_y)$ with a set of random hidden variables \mathbf{h} , it is difficult to obtain a close form of $P(\mathbf{x}|\theta_y)$ since $\int P(\mathbf{x}, \mathbf{h}|\theta_y) d\mathbf{h}$ is usually intractable. We can resort to the following variational lower bound [7], [4]:

$$\log P(\mathbf{x}|\theta_y) \geq E_{Q(\mathbf{h})}[\log P(\mathbf{x}, \mathbf{h}) - \log Q(\mathbf{h})] \triangleq F(\mathbf{x}, \theta_y) \quad (2)$$

where $Q(\mathbf{h})$ is the variational approximate posterior of $P(\mathbf{h}|\mathbf{x})$. Then, instead of the intractable discriminant function (Eq. (1)), we resort to the following tractable one [19], [9]

$$\hat{L}(\mathbf{x}, \Theta) = F(\mathbf{x}, \theta_+) - F(\mathbf{x}, \theta_-) + b \quad (3)$$

We assume the generative model $P(\mathbf{x}, \mathbf{h}|\theta)$ belong to the exponential family which covers most models. We have the general form $P(\mathbf{x}, \mathbf{h}) = \exp\{a(\theta)^T T(\mathbf{x}, \mathbf{h}) + S(\mathbf{x}, \mathbf{h}) + d(\theta)\}$ where θ is the vector of parameters; $T(\mathbf{x}, \mathbf{h})$ is the vector of sufficient statistics; $S(\mathbf{x}, \mathbf{h})$ and $d(\theta)$ are scalar functions. Similarly, the prior over \mathbf{h} is $P(\mathbf{h}) = \exp\{c(\theta_h)^T T(\mathbf{h}) + S(\mathbf{h}) + f(\theta_h)\}$. Further, we assume that the approximate posterior of \mathbf{h} , for the example \mathbf{x} , takes the same form with its prior $P(\mathbf{h})$ but with different parameter [4] $Q(\mathbf{h}) = \exp\{c(\theta'_h)^T T(\mathbf{h}) + S(\mathbf{h}) + f(\theta'_h)\}$. Substituting the above formulas of $P(\mathbf{x}, \mathbf{h})$ and $Q(\mathbf{h})$ into Eq. (2), it can be verified that $F(\mathbf{x}, \theta) = E_{Q(\mathbf{h})}[\log P(\mathbf{x}, \mathbf{h}) - \log Q(\mathbf{h})] = \alpha^T E_{Q(\mathbf{h})}[\tilde{T}(\mathbf{x}, \mathbf{h})] + \beta$, where $\alpha = (a(\theta)^T, 1, -1^T, -1)$; $\beta = d(\theta)$; $\tilde{T}(\mathbf{x}, \mathbf{h}) = (T(\mathbf{x}, \mathbf{h})^T, S(\mathbf{x}, \mathbf{h}), (\text{diag}(c(\theta'_h)^T T(\mathbf{h}))^T, S(\mathbf{h}), f(\theta'_h))^T$. For a pair of models θ_+ and θ_- , Eq. (3) can be written as:

$$\begin{aligned} \hat{L}(\mathbf{x}, \Theta) &= \alpha_+^T E_{Q(\mathbf{h}_+)}[\tilde{T}(\mathbf{x}, \mathbf{h}_+)] - \alpha_-^T E_{Q(\mathbf{h}_-)}[\tilde{T}(\mathbf{x}, \mathbf{h}_-)] + \beta_+ - \beta_- + b \\ &= \tilde{\alpha}^T E_{Q(\mathbf{h}_+, \mathbf{h}_-)}[\phi(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-)] \end{aligned} \quad (4)$$

where $\tilde{\alpha} = (\alpha_+^T, -\alpha_-^T, \beta_+ - \beta_- + b)^T$, $\phi(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-) \triangleq (\tilde{T}(\mathbf{x}, \mathbf{h}_+)^T, \tilde{T}(\mathbf{x}, \mathbf{h}_-)^T, 1)^T$. Eq. (4) takes the form of the linear classifier, where $E_Q[\phi(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-)]$ is considered to furnish a feature mapping. From another perspective, $\phi(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-)$ can be considered as a *stochastic feature mapping* because the hidden variables \mathbf{h}_+ , \mathbf{h}_- are all conditioned on the example \mathbf{x} and thus its value can serve as feature for identifying \mathbf{x} . It is considered to define a *stochastic feature space* because it is evaluated based on stochastic examples drawn from the posterior of \mathbf{h} .

III. PAC-BAYES BOUND STOCHASTIC CLASSIFIER

The derived stochastic feature mapping $\phi(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-)$ makes it possible to jointly learn generative models (subsequently

feature mappings) and classifier. We construct a linear Gibbs classifier over this stochastic feature mapping:

$$G_Q = \text{sign}[\mathbf{w} \cdot \phi(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-)] \triangleq f_{\mathbf{w}}(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-) \quad (5)$$

where \mathbf{w} is the weight of classifier; \mathbf{h}_+ , \mathbf{h}_- , \mathbf{w} follow some distribution Q which will be specified in Section III-B. Gibbs classifier with such a feature mapping offers *several advantages*. *First*, this classifier allows PAC-Bayes risk bounds that have explicit solutions for $Q(\mathbf{h}_+, \mathbf{h}_-)$ which can help tune the feature mapping for better classification; *Second*, the PAC-Bayes risk bound for such a classifier can be tighter than VC bounds [11]; *Third*, the feature mapping is independent with model parameters θ , making the solution of θ very simple.

A. PAC-Bayes bounds for stochastic feature mapping

Let \mathcal{X} be the input space consisting of an arbitrary subset of \mathbb{R}^d and $\mathcal{Y} = \{-1, +1\}$ be the output space. An example is an input-output pair (\mathbf{x}, y) where $\mathbf{x} \in \mathcal{X}$ and $y \in \mathcal{Y}$. In a PAC-Bayes setting [13], each example (\mathbf{x}, y) is drawn from a fixed, but unknown, probability distribution D on $\mathcal{X} \times \mathcal{Y}$. Let $f(\mathbf{x}, \mathbf{v}) : \mathcal{X} \rightarrow \mathcal{Y}$ be any classifier with a set of variables \mathbf{v} . The learning task is to choose a posterior distribution Q over a space \mathcal{F} of classifiers and a space \mathcal{V} of variables such that the Q -weight majority classifier $B_Q = \text{sign}[E_{(f, \mathbf{v}) \sim Q} f(\mathbf{x}, \mathbf{v})]$ will have the smallest possible risk on the training example set $S = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m)\}$. The output of B_Q is closely related to the output of the Gibbs classifier G_Q which first chooses a classifier f and a vector \mathbf{v} according to Q , and then classifies an example \mathbf{x} . The *true risk* $R(G_Q)$ and the *empirical risk* $R_S(G_Q)$ of this Gibbs classifier are given by:

$$R(G_Q) = E_{(f, \mathbf{v}) \sim Q} E_{(\mathbf{x}, y) \sim D} \mathbb{I}(f(\mathbf{x}, \mathbf{v}) \neq y) \quad (6)$$

$$R_S(G_Q) = E_{(f, \mathbf{v}) \sim Q} \frac{1}{m} \sum_{i=1}^m \mathbb{I}(f(\mathbf{x}_i, \mathbf{v}) \neq y_i) \quad (7)$$

This setting is naturally accommodated by PAC-Bayes theory since \mathbf{v} can be considered as a part of f . Among several PAC-Bayes bounds [13], [5], [2], [12], the bound derived in [2] is quite tight and gives an explicit bound for the true risk $R(G_Q)$, which allows the derivation of the posterior Q , in contrast to most of the other implicit bounds over $\text{KL}(R_S(G_Q) \| R(G_Q))$.

Theorem 1. *For any distribution D over $\mathcal{X} \times \mathcal{Y}$, any space \mathcal{F} of classifiers, any space \mathcal{V} of random variables, any distribution P over $\mathcal{F} \times \mathcal{V}$, any $\delta \in (0, 1]$, and any real number $C > 0$, $\forall Q$ over $\mathcal{F} \times \mathcal{V}$, we have:*

$$\Pr \left(R(G_Q) \leq \frac{1}{1 - e^{-C}} \left[1 - \exp \left\{ -CR_S(G_Q) - \frac{1}{m} [\text{KL}(Q \| P) - \ln \delta] \right\} \right] \right) \geq 1 - \delta$$

This is a slight extension of Corollary 2.2 of [2], and can be proved by replacing f with (f, \mathbf{v}) and reapplying its proof [2].

The above risk bound is derived for labeled data. Here we have extended the bound to accommodate both labeled and unlabeled data for semi-supervised learning in the following theorem. The semi-supervised bound is different with [12] whose bound is implicit and has no explicit solution of Q .

Theorem 2. For any distribution D over $\mathcal{X} \times \mathcal{Y}$, any space \mathcal{F} of classifiers, any space \mathcal{V} of random variables, distribution P over $\mathcal{F} \times \mathcal{V}$, any $\delta \in (0, 1]$, and any real number $C > 0$, $\forall Q$ over $\mathcal{F} \times \mathcal{V}$, we have:

$$\Pr \left(R(G_Q) \leq \frac{1}{1-e^{-C}} \left[1 - \exp \left\{ -C \left[e_S(G_Q) + \frac{1}{2} d_S(G_Q) \right] - \frac{1}{m} [\text{KL}(Q \| P) - \ln \delta] \right\} \right] \right) \geq 1 - \delta$$

where the risks for labeled and unlabeled data are $e_S(G_Q) = E_{(f_1, \mathbf{v}_1) \sim Q} E_{(f_2, \mathbf{v}_2) \sim Q} \frac{1}{m} \sum_{i=1}^m \mathbf{I}(f_1(\mathbf{x}_i, \mathbf{v}_1) \neq y_i) \mathbf{I}(f_2(\mathbf{x}_i, \mathbf{v}_2) \neq y_i)$ and $d_S(G_Q) = E_{(f_1, \mathbf{v}_1) \sim Q} E_{(f_2, \mathbf{v}_2) \sim Q} \frac{1}{m} \sum_{i=1}^m \mathbf{I}(f_1(\mathbf{x}_i, \mathbf{v}_1) \neq f_2(\mathbf{x}_i, \mathbf{v}_2))$.

Proof: Let E_f be the abbreviation of $E_{(f, \mathbf{v}) \sim Q}$. Note that $E_{f_1, f_2} \mathbf{I}(f_1 \neq f_2) = E_{f_1, f_2} 2\mathbf{I}(f_1 \neq y) \mathbf{I}(f_2 = y) = E_{f_1, f_2} 2\mathbf{I}(f_1 \neq y)(1 - \mathbf{I}(f_2 \neq y)) = E_{f_1, f_2} 2(\mathbf{I}(f_1 \neq y) - \mathbf{I}(f_1 \neq y) \mathbf{I}(f_2 \neq y))$ and $R_S(G_Q) = \frac{1}{m} \sum_i E_{f_1} \mathbf{I}(f_1 \neq y_i)$ (Eq. (7)). Therefore $R_S(G_Q) = \frac{1}{m} \sum_i E_{f_1} \mathbf{I}(f_1^i \neq y_i) = \frac{1}{m} \sum_i E_{f_1, f_2} \mathbf{I}(f_1^i \neq y_i) \mathbf{I}(f_2 \neq y_i) + \frac{1}{2m} \sum_i E_{f_1, f_2} \mathbf{I}(f_1^i \neq f_2^i) = e_S(G_Q) + \frac{1}{2} d_S(G_Q)$, ($f_1^i = f_1(\mathbf{x}_i)$). Substituting $R_S(G_Q) = e_S(G_Q) + \frac{1}{2} d_S(G_Q)$ into Theorem 1, then we obtain Theorem 2. ■

Since $d_S(G_Q)$ is independent of labels, it allows classifiers using the above bound to exploit unlabeled data. Minimizing this risk $d_S(G_Q)$ would contract the posteriors over the stochastic classifier and the stochastic feature space, making classification and feature mapping less uncertain.

B. Objective function and specification

Let $B(Q, C) = \frac{1}{1-e^{-C}} [1 - \exp\{-J(Q) + \frac{1}{m} \ln \delta\}]$ be the upper bound in Theorem 2, where $J(Q) = C(e_S(G_Q) + \frac{1}{2} d_S(G_Q)) + \frac{1}{m} \text{KL}(Q \| P)$. Training a classifier with minimum risk means minimizing the upper bound $B(Q, C)$ w.r.t. Q and C . Note that minimizing $B(Q, C)$ w.r.t. Q equals to minimizing $J(Q)$ w.r.t. Q . Since unlabeled data are only available in estimating $d_S(G_Q)$, $J(Q)$ over labeled data S_l of m_l examples and unlabeled data S_u of m_u examples can be written as

$$J(Q) = C \left[e_{S_l}(G_Q) + \frac{1}{2} d_{S_u}(G_Q) \right] + \frac{1}{m} \text{KL}(Q \| P)$$

where $m = m_l + m_u$. This form enables us to derive the analytical form of posterior distributions Q of the classifier and the hidden variables. Apply the above bound to the stochastic classifier defined in Eq. (5) and set $\mathbf{v} = (\mathbf{h}_+, \mathbf{h}_-)$, we have

$$f = f_{\mathbf{w}}, \quad E_{f \sim Q}[\cdot] = E_{\mathbf{w} \sim Q}[\cdot] \quad (8)$$

Then learning PAC-Bayes bound classifier with the generative model embedding is to minimize the objective function $J(Q)$ w.r.t. the posterior $Q(\mathbf{w}, \mathbf{h}_+, \mathbf{h}_-)$ and parameters θ_+, θ_- :

$$J(Q) = C \left[e_{S_l}(G_Q) + \frac{1}{2} d_{S_u}(G_Q) \right] + \frac{1}{m} [\text{KL}(Q(\mathbf{w}) \| P(\mathbf{w})) + \text{KL}(Q(\mathbf{h}_+) \| P(\mathbf{x}, \mathbf{h}_+ | \theta_+)) + \text{KL}(Q(\mathbf{h}_-) \| P(\mathbf{x}, \mathbf{h}_- | \theta_-))] \quad (9)$$

where G_Q is the linear stochastic classifier defined in Eq. (5); $P(\mathbf{x}, \mathbf{h}_+ | \theta_+)$ and $P(\mathbf{x}, \mathbf{h}_- | \theta_-)$ are generative models for positive and negative classes respectively; the first row is the objective function for regular Gibbs classifier and the second row is the objective function for two generative models.

To compute the objective function $J(Q)$ in Eq. (9), we will need to have approximations or expressions for e_{S_l} , d_{S_u} and

$\text{KL}(Q \| P)$ that are computationally tractable. To derive these expressions, as were done in [5], we assume that the prior of the weight is Gaussian $P(\mathbf{w}) = N(\mathbf{u}_0, \mathbf{I})$ and its posterior is also Gaussian except with a different mean, i.e., $Q(\mathbf{w}) = N(\mathbf{u}, \mathbf{I})$. Based on this assumption, we have:

$$\text{KL}(Q(\mathbf{w}) \| P(\mathbf{w})) = \frac{1}{2} \|\mathbf{u} - \mathbf{u}_0\|^2 \quad (10)$$

Using the assumption and Gaussian integrals [5], we have,

$$E_{\mathbf{w} \sim Q} \mathbf{I}(f_{\mathbf{w}}(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-) \neq y) = \Phi(y\mathbf{u} \cdot \bar{\phi}(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-)) \quad (11)$$

where $\Phi(a) = \int_a^\infty \frac{1}{\sqrt{2\pi}} \exp(-\frac{x^2}{2}) dx = \frac{1}{2} \text{erfc}(\frac{a}{\sqrt{2}})$ and $\bar{\phi} = \frac{\phi(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-)}{\|\phi(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-)\|}$. Further, considering Eq. (11), we have the integration:

$$\begin{aligned} E_{\mathbf{w}_1, \mathbf{w}_2 \sim Q} \mathbf{I}(f_{\mathbf{w}_1} \neq f_{\mathbf{w}_2}) &= E_{\mathbf{w}_1, \mathbf{w}_2 \sim Q} 2\mathbf{I}(f_{\mathbf{w}_1} \neq 1) \mathbf{I}(f_{\mathbf{w}_2} \neq -1) \\ &= 2\Phi(\mathbf{u} \cdot \bar{\phi}(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-)) \Phi(-\mathbf{u} \cdot \bar{\phi}(\mathbf{x}, \mathbf{h}_+, \mathbf{h}_-)) \end{aligned} \quad (12)$$

With these formulas, we proceed to obtain an expression for $J(Q)$, and find $Q(\mathbf{h}_+, \mathbf{h}_-)$, $Q(\mathbf{w})$, θ_+ and θ_- by minimizing $J(Q)$ with an iterative optimization procedure in the next section.

IV. INFERENCE AND PARAMETER ESTIMATION

In this section, we derive the learning procedure (inference and parameter estimation) of the proposed approach. Consider Eq. (9) and $e_S(G_Q) = \frac{1}{m} \sum_i E_{\mathbf{w}_1} \mathbf{I}(f_{\mathbf{w}_1}(\mathbf{x}_i) \neq y_i) E_{\mathbf{w}_2} \mathbf{I}(f_{\mathbf{w}_2}(\mathbf{x}_i) \neq y_i)$ and $d_S(G_Q) = \frac{1}{m} \sum_i E_{\mathbf{w}_1, \mathbf{w}_2} \mathbf{I}(f_{\mathbf{w}_1}(\mathbf{x}_i) \neq f_{\mathbf{w}_2}(\mathbf{x}_i))$ (Th.2), $J(Q)$ (average cost) over the training set $S = S_l \cup S_u$ is:

$$\begin{aligned} J(Q) &= C \left[\frac{1}{m_l} \sum_{i=1}^{m_l} E_Q \mathbf{I}(f_{\mathbf{w}}(\mathbf{x}_i) \neq y_i) + \frac{1}{m_u} \sum_{i=1}^{m_u} E_Q \mathbf{I}(f_{\mathbf{w}_1}(\mathbf{x}_i) \neq f_{\mathbf{w}_2}(\mathbf{x}_i)) \right] \\ &\quad + \frac{1}{m^2} \sum_{i=1}^m \text{KL}(Q_i(\mathbf{h}_+, \mathbf{h}_-) \| P(\mathbf{x}_i, \mathbf{h}_+, \mathbf{h}_-)) + \frac{1}{m} \text{KL}(Q(\mathbf{w}) \| P(\mathbf{w})) \end{aligned}$$

where the terms $E_Q \mathbf{I}(f_{\mathbf{w}}(\mathbf{x}_i) \neq y_i)$, $E_Q \mathbf{I}(f_{\mathbf{w}_1}(\mathbf{x}_i) \neq f_{\mathbf{w}_2}(\mathbf{x}_i))$ and $\text{KL}(Q(\mathbf{w}) \| P(\mathbf{w}))$ have been respectively given by Eq. (11), Eq. (12) and Eq. (10); $P(\mathbf{x}_i, \mathbf{h}_+, \mathbf{h}_-)$ is the abbreviation of $P(\mathbf{x}_i, \mathbf{h}_+)P(\mathbf{x}_i, \mathbf{h}_-)$. We now show how an EM-like iterative procedure [4] can be used to learn the stochastic feature space and the Gibbs classifier simultaneously.

A. Inference: minimize $J(Q)$ w.r.t. $Q_i(\mathbf{h}_+, \mathbf{h}_-)$

In the *first step*, we fix $Q(\mathbf{w})$, θ_+ , θ_- , and minimize $J(Q)$ w.r.t. $Q_i(\mathbf{h}_+, \mathbf{h}_-)$, subject to $\int Q_i(\mathbf{h}_+, \mathbf{h}_-) d\mathbf{h}_+ d\mathbf{h}_- = 1$. Benefiting from the explicit bound (Th. 2), we have the following solution:

$$Q_i(\mathbf{h}_+, \mathbf{h}_-) = \frac{1}{Z_i} P(\mathbf{h}_+, \mathbf{h}_-, \mathbf{x}_i) \exp\{-C E_{Q(\mathbf{w})}[\varphi_i]\} \quad (13)$$

where $\varphi_i = \frac{m^2}{m_l} \mathbf{I}(f_{\mathbf{w}} \neq y_i)$ if $\mathbf{x}_i \in S_l$ and $\varphi_i = \frac{m^2}{2m_u} \mathbf{I}(f_{\mathbf{w}_1} \neq f_{\mathbf{w}_2})$ if $\mathbf{x}_i \in S_u$. Note that $E_{Q(\mathbf{w})}[\varphi_i]$ is given by Eq. (11) and Eq. (12). The fact that the output of classifier is inside the expression for posteriors means that the generative models are being tuned as well when the classifier is being optimized during the minimization of PAC-Bayes bound. This tuning inhibits those examples of \mathbf{h}_+ , \mathbf{h}_- that lead to misclassification and encourages those with less misclassification. Sampling from this posterior is simple using *Gibbs-rejection sampling*, because $P(\mathbf{h}, \mathbf{x}_i)$ can be directly used as the comparison function

since $P(\mathbf{h}, \mathbf{x}_i) \exp(\cdot) \leq P(\mathbf{h}, \mathbf{x}_i)$ ($E_{Q(\mathbf{w})}[\varphi_i] \geq 0$ as $I(\cdot)$ is a zero-one output function). Considering the j -th example \mathbf{h}_{ij} drawn from $P(\mathbf{h}, \mathbf{x}_i)$, we reject it if $Q_i(\mathbf{h}_{ij}) < r_j$ where r_j is drawn from the uniform distribution over $[0, P(\mathbf{h}_{ij}, \mathbf{x}_i)]$.

B. Parameter estimation: minimize $J(Q)$ w.r.t parameters

In the *second step*, we fix the posteriors $Q_i(\mathbf{h}_+, \mathbf{h}_-)$ and parameters θ_+, θ_- , and determine the posterior distribution $Q(\mathbf{w})$. Instead of sampling from $Q(\mathbf{w})$, we directly determine its parameter \mathbf{u} by minimizing $J(Q)$ w.r.t. \mathbf{u} . Since $J(Q)$ w.r.t. \mathbf{u} is intractable, we resort to minimizing its upper bound $J(\mathbf{u})$ (see the Appendix) w.r.t. \mathbf{u} . The gradient of $J(\mathbf{u})$ w.r.t. \mathbf{u} is:

$$\frac{\partial J(\mathbf{u})}{\partial \mathbf{u}} = \frac{1}{m}(\mathbf{u} - \mathbf{u}_0) - \frac{C}{mIn} \sum_{i,j=1}^{m_i, n} g(y_i \mathbf{u} \cdot \bar{\phi}_{ij}) y_i \bar{\phi}_{ij} \quad (14)$$

$$+ \frac{C}{m_u n} \sum_{i,j=1}^{m_u, n} g(\mathbf{u} \cdot \bar{\phi}_{ij}) [\Phi(\mathbf{u} \cdot \bar{\phi}_{ij}) - \Phi(-\mathbf{u} \cdot \bar{\phi}_{ij})] \bar{\phi}_{ij}$$

where $g(\cdot)$ is the gaussian function with mean $\mu = 0$ and std $\delta = 1$. The gradient of $B(Q, C)$ with respect to C is:

$$\frac{\partial B(Q, C)}{\partial C} = \frac{-e^{-C}}{(1 - e^{-C})^2} \left[1 - \exp\left(-J(Q) - \frac{1}{m} \log \delta\right) \right] \quad (15)$$

$$+ \frac{1}{1 - e^{-C}} [R_{S_l}(G_Q) + d_{S_u}(G_Q)] \exp\left(-J(Q) - \frac{1}{m} \log \delta\right)$$

In the *third step*, we fix $Q_i(\mathbf{h}_+, \mathbf{h}_-)$, \mathbf{u} and update parameters θ_+, θ_- . Note only the third term of Eq. (9), i.e., the objection function of the positive model, involves θ_+ . So the update rules of θ_+ , derived by minimizing Eq. (9) w.r.t. θ_+ , are same as those of the original generative model. Similarly for θ_- .

The *learning* procedure is summarized in Algorithm 1. In *classification*, similar with [2], we use the decision rule of majority vote $\hat{y} = \text{sign}[\frac{1}{n} \sum_{j=1}^n E_{P(\mathbf{h}_+, \mathbf{h}_- | \mathbf{x}_i) Q(\mathbf{w})} \mathbf{w} \cdot \phi(\mathbf{x}_i, \mathbf{h}_+, \mathbf{h}_-)] \approx \text{sign}[\frac{1}{n} \sum_{j=1}^n \mathbf{u} \cdot \phi(\mathbf{x}_i, \mathbf{h}_{+ij}, \mathbf{h}_{-ij})]$ with $n = 5$ and $(\mathbf{h}_{+ij}, \mathbf{h}_{-ij})$ being the j -th example drawn from $P(\mathbf{h}_+, \mathbf{h}_- | \mathbf{x}_i)$.

Algorithm 1 Inference and learning

- 1: **input:** data set S_l, S_u , and S'_l, S'_u are fractions of S_l, S_u
 - 2: initialize $\hat{\mathbf{u}}, \hat{\theta}_+, \hat{\theta}_-$, learning rates γ_u, γ_c , and $\delta = 0.05$
 - 3: $\hat{\mathbf{u}}_0 \leftarrow \min_{\mathbf{u}} R_{S'_l}(G_Q) + \frac{1}{2} d_{S'_u}(G_Q)$
 - 4: **repeat**
 - 5: **for** $i = 1$ to m **do**
 - 6: sample $Q_i(\mathbf{h}_+, \mathbf{h}_-)$ using Gibbs-rejection sampling
 - 7: **end for**
 - 8: update $\hat{\theta}_+$ with $\{\mathbf{h}_{+ij}\}_{ij} (\mathbf{x}_i \in S_l^+)$ using the rules of the original generative model. Similar for $\hat{\theta}_-$.
 - 9: $\hat{\mathbf{u}} \leftarrow \hat{\mathbf{u}} - \gamma_u \frac{\partial J(\mathbf{u})}{\partial \mathbf{u}}, \quad C \leftarrow C - \gamma_c \frac{\partial B}{\partial C}$
 - 10: **until** convergence
 - 11: **output:** $\hat{\mathbf{u}}_0, \hat{\mathbf{u}}, \hat{\theta}_+, \hat{\theta}_-$
-

V. EXPERIMENTS

This section empirically evaluates the proposed method stochastic feature mapping (SFM) and related methods on general classification tasks, scene recognition and protein sequence classification respectively. For multiple-class classification problem, we divide it into binary classification problems, each of which is an *one-versus-rest problem* that distinguishes




Fig. 1: Classification accuracy (%) for varying number of training examples. left: Tissue (UCI); middle: Highway (scene); right: Superfamily #2 (protein).

one class from others. For each binary problem, we randomly partition the positive examples into 50% training and 50% test sets, and similarly for negative examples. We test each binary classification problem on 20 random partitions, and report the average results. For the semi-supervised version, we use 25% of test examples as unlabeled data. Two *related and general methods*, Fisher score (FS) [3] and free energy score space (FESS) [9], and some other state-of-the-art methods are also tested for comparison.

There are two points in implementation. First, the optimization procedures of \mathbf{u}_0 and \mathbf{u} may suffer from the local minima problem, resulting in poor solution. The strategy adopted by [2] is to perform the optimization for 10~100 trials where a new random initial point within the range $[-20, 20]^d$ is used in each trial. Second, the value of parameter C has been shown to be important. Another effective strategy experimented is to assess the performance using 10-fold cross-validation.

A. Deriving a general classification tool

In the first experiment, we derive a general classification method by applying the proposed framework to a simple yet general generative model, Gaussian mixture model. Let $\mathbf{x} \in \mathbb{R}^d$ be the observed variable; $\mathbf{z} = \{z_1, \dots, z_k\}$ be the hidden binary indicate vector for K mixture components, and assume the covariance matrix be diagonal; $\mathbf{a} = \{a_1, \dots, a_k\}$ be the parameters of the approximate posterior of \mathbf{z} . The elements of the feature mapping ϕ of this model are $\{z_i(\mathbf{x}^T, \text{diag}(\mathbf{xx}^T), 1), z_i \log a_i\}_{i=1}^K$. The posterior of \mathbf{z} can be easily derived from Eq. (13). The number of mixture components is configured to $K = 4$ throughout the experiment.

We select 8 data sets from UCI database for evaluation, preferring those with no missing entities. The number of classes of each data set is between 2 and 15. The number of examples of each class varies from 14 and 673. The dimensionality is between 9 and 90. We compare our method SFM with Adaboost [17], SVM [11], localized multiple kernel learning (LMKL) [18] and PAC-Bayes gradient descent PBGD3 [2]. The average results are reported in Table I. It shows that SFM is adaptive to different data sets and outperforms other methods in half of the data sets. It is also worth noting that the linear version of PBGD3 does work well in these evaluation. The results of semi-supervised version is presented in Fig. 1.

B. Scene recognition

We evaluate our SFM method and compare its performance against comparable methods on a typical vision task, scene recognition. In this task, visual words are used for image representation for its robustness to topic and spatial variance.

TABLE I: Classification accuracy (% \pm std) on UCI database (*one-versus-rest* on each dataset).

DATA	Adaboost	SVM	LMKL	PBGD3	SFM-GMM
Cancer	93.24 \pm 1.26	96.80 \pm 1.79	96.41 \pm 0.97	93.98 \pm 1.52	95.02 \pm 0.95
Tissue	88.55 \pm 5.91	78.17 \pm 12.27	87.69 \pm 5.24	88.14 \pm 0.50	89.04 \pm 3.12
Wine	92.98 \pm 3.42	97.73 \pm 1.86	95.48 \pm 4.10	92.22 \pm 12.63	95.79 \pm 1.33
Sonar	70.87 \pm 4.76	73.11 \pm 3.25	80.21 \pm 1.52	75.52 \pm 5.70	80.60 \pm 4.24
Credit	84.74 \pm 1.30	84.74 \pm 1.79	81.92 \pm 1.41	83.53 \pm 1.82	84.89 \pm 1.30
SpHeart	78.65 \pm 2.08	74.66 \pm 3.56	80.38 \pm 3.40	79.70 \pm 0.65	80.84 \pm 0.49
Libras	92.97 \pm 1.76	87.54 \pm 7.01	96.58 \pm 1.78	94.52 \pm 2.80	95.61 \pm 3.34
Steel	89.47 \pm 9.08	86.43 \pm 9.16	92.63 \pm 8.14	87.30 \pm 8.26	89.10 \pm 8.49

TABLE II: Accuracy (% \pm std.) of *one-versus-rest* scene recognition.

SCENE	PHOW-SVM	LDA-MAP	FS-LDA	FESS-LDA	SFM-LDA
Coast	90.66 \pm 0.65	83.85 \pm 0.92	90.42 \pm 0.34	93.89 \pm 0.46	94.06 \pm 0.64
Forest	96.49 \pm 0.39	94.94 \pm 0.46	94.45 \pm 0.46	97.92 \pm 0.26	98.10 \pm 0.32
Mountain	92.58 \pm 0.64	84.99 \pm 1.78	88.62 \pm 0.50	93.29 \pm 0.47	93.80 \pm 0.44
Country	91.38 \pm 0.71	72.30 \pm 1.74	87.40 \pm 0.46	90.62 \pm 0.33	90.82 \pm 0.62
Highway	95.27 \pm 0.49	81.50 \pm 1.28	92.48 \pm 0.22	94.67 \pm 0.34	95.46 \pm 0.30
InsideCity	93.96 \pm 0.62	85.14 \pm 1.74	90.79 \pm 0.14	94.26 \pm 0.65	95.27 \pm 0.35
Street	93.89 \pm 0.64	76.46 \pm 1.23	93.76 \pm 0.24	94.21 \pm 0.42	95.03 \pm 0.47
Building	94.40 \pm 0.49	87.85 \pm 0.55	92.83 \pm 0.57	96.06 \pm 0.51	95.81 \pm 0.38

We use latent Dirichlet allocation (LDA) [1] to model the distributions of visual words, and derive a recognition tool under the proposed framework. Like [15], we sample the topic variable using collapsed Gibbs sampling and reject examples according to the rule for Eq. (13). We fix the parameter α and allow β [15] to be updated. Let w, z respectively indicate word and topic, and γ be the parameter of the approximate posterior of z . The elements of the feature mapping ϕ of a model are $\{z_{nk}, w_n z_{nk}, z_{nk} \log \gamma_{nk} \}_{n,i,k}$ where n, i, k index word, term and topic respectively. For FS [3] and FESS [9], we extract features from the trained LDA model and deliver to SVM. The number of topics of LDA is set to 50.

The CVCL scene dataset is chosen for evaluation. It contains 4 artificial scenes and 4 natural scenes. For each image, dense SIFT descriptors [6] are extracted from 20×20 grid patches over 4 scales. These descriptors are quantized to visual words using a code book (50 centers) clustered from some random selected descriptors. The resulting visual words of an image are in the form of histogram where each bin corresponds to a code center of the code book. The evaluation results are summarized in Table II. Our results compare well with PHOW [16] which is a state-of-the-art feature for scene recognition. The results of semi-supervised learning are shown in Fig. 1, demonstrating unlabeled examples can help classification particularly when there are few labeled examples.

C. Protein classification

To evaluate the capability of the proposed approach in dealing with variable length sequences, we apply the proposed framework to remote homology recognition. The problem here assigns test protein sequences to the domain superfamilies

defined in the SCOP (1.53) taxonomy tree according to functions of proteins. The protein sequence data is obtained from ASTRAL database with E-value threshold of 10^{-25} to reduce similar sequences. We uses four labeled domain superfamilies, metabolism, information, intra-cellular processes and extra-cellular processes for evaluation. The numbers of sequences are 804, 950, 695 and 992 respectively. Each protein sequence is a string composed of 22 distinct letters, and the string length varies from 20 to 994.

Hidden Markov model (HMM) [10] is used to model the distribution over protein sequences for its ability in handling sequences with variable length. The number of output states is 22, and the number of hidden states is set to 10. Let \mathbf{x} be the sequence with length T_x , where \mathbf{x}^t be the binary indicator where $x_k^t = 1$ if the k -th state of K possible ones is selected at time t . Let \mathbf{q}^t be the binary state indicator where $q_i^t = 1$ if the i -th state of M possible ones is selected at time t ; $A_{M \times M}$ be the transition probabilities of the approximate posteriors. The elements of the feature mapping ϕ can be written as $\{q_i^0, \sum_{t=0}^{T_x-1} q_i^t q_j^{t+1}, \sum_{t=0}^{T_x-1} q_i^t q_j^{t+1} \log A_{ij}, \sum_{t=0}^{T_x} q_i^t x_k^t\}_{i,j,k}$. With the hidden states of the input sequence inferred by Baum-Welch algorithm [22], it is easy to estimate the posterior transition probabilities conditioned on \mathbf{x} . Using the sampling distribution derived in Eq. (13), we are able to draw the examples of hidden states and re-estimate their posterior. The results are reported in Table III. The 2-gram feature is actually the transition probability of observed states of a sequence, i.e. $\{\frac{1}{T_x} \sum_{t=0}^{T_x-1} x_i^t x_k^{t+1}\}_{i,k}$. The difference of the performance of the first four methods are not significant except on family #3. The results of semi-supervised learning are reported in Fig. 1, which shows improvement on few training samples.

TABLE III: Accuracy (%±std.) of one-versus-rest protein recognition.

SUP.FAM.	2GRAM-SVM	HMM-MAP	FS-HMM	FESS-HMM	SFM-HMM
# 1	78.79 ± 1.13	80.91 ± 1.53	80.03 ± 0.78	80.12 ± 0.84	82.75 ± 0.96
# 2	79.01 ± 0.97	80.10 ± 0.51	77.56 ± 0.64	78.96 ± 0.59	83.08 ± 0.55
# 3	75.19 ± 0.86	77.92 ± 0.79	73.31 ± 0.21	73.35 ± 0.41	79.28 ± 0.64
# 4	96.01 ± 0.33	95.10 ± 0.39	94.27 ± 0.37	97.58 ± 0.13	96.72 ± 0.63

VI. CONCLUSIONS

This paper presents a framework to incorporate the abilities of generative model and discriminative model for classification under the PAC-Bayes theory. The bridge of this incorporation is a stochastic feature mapping which is derived from the linear form of the practical MAP classifier and is independent with the parameters of the adopted generative models. Under this framework, the derived stochastic feature mapping and generative models can be tuned during the training of the classifier. A major difficulty is the non-convexity of the objective function, where local minima can hamper the solution. Our approach can benefit from the development or exploitation of more robust and efficient optimization methods.

APPENDIX

Since $J(Q)$ is intractable, we derive its upper bound by fixing θ_+ , θ_- , $Q_i(\mathbf{h}_+, \mathbf{h}_-)$. Using Eq. (10) and Eq. (13), we have

$$\begin{aligned}
\text{KL}(Q \| P) &= \text{KL}(Q(\mathbf{w}, \mathbf{h}_+, \mathbf{h}_-) \| P(\mathbf{w})P(\mathbf{x}, \mathbf{h}_+ | \theta_+)P(\mathbf{x}, \mathbf{h}_- | \theta_-)) \\
&= \mathbb{E}_{Q_w} \left[\log \frac{Q(\mathbf{w})}{P(\mathbf{w})} \right] + \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{Q_i} \left[\log \frac{Q_i(\mathbf{h}_+)Q_i(\mathbf{h}_-)}{P(\mathbf{x}_i, \mathbf{h}_+)P(\mathbf{x}_i, \mathbf{h}_-)} \right] \\
&= \frac{1}{2} \|\mathbf{u} - \mathbf{u}_0\|^2 - Cm \left[e_{S_l}(G_Q) + \frac{1}{2} d_{S_u}(G_Q) \right] - \frac{1}{m} \sum_{i=1}^m \log Z_i
\end{aligned}$$

where $Q_w = Q(\mathbf{w})$, $Q_i = Q_i(\mathbf{h}_+, \mathbf{h}_-)$ and

$$\begin{aligned}
\sum_{i=1}^m \log Z_i &= \sum_{i=1}^m \log \mathbb{E}_{Q_i} [\exp \{-C \mathbb{E}_{Q_w}[\varphi_i]\}] \geq \sum_{i=1}^m \mathbb{E}_{Q_i} [-C \mathbb{E}_{Q_w}[\varphi_i]] \\
&= \sum_{i=1}^{m_l} \mathbb{E}_{Q_w Q_i} \left[\frac{Cm^2}{m_l} \mathbb{I}(f_w^i \neq y_i) \right] + \sum_{i=1}^{m_u} \mathbb{E}_{Q_w Q_i} \left[\frac{Cm^2}{2m_u} \mathbb{I}(f_{w_1}^i \neq f_{w_2}^i) \right] \\
&\approx \frac{Cm^2}{m_l n} \sum_{i,j=1}^{m_l n} \Phi(y_i \mathbf{u} \cdot \bar{\phi}_{ij}) + \frac{Cm^2}{m_u n} \sum_{i,j=1}^{m_u n} \Phi(\mathbf{u} \cdot \bar{\phi}_{ij}) \Phi(-\mathbf{u} \cdot \bar{\phi}_{ij})
\end{aligned}$$

where the inequality is derived by applying Jensen's inequality; $f_w^i = f_w(\mathbf{x}_i)$; $\bar{\phi}_{ij} = \frac{\phi(\mathbf{x}_i, \mathbf{h}_{+ij}, \mathbf{h}_{-ij})}{\|\phi(\mathbf{x}_i, \mathbf{h}_{+ij}, \mathbf{h}_{-ij})\|}$ where $(\mathbf{h}_{+ij}, \mathbf{h}_{-ij})$ represents the j -th example drawn from $Q_i(\mathbf{h}_+, \mathbf{h}_-)$. Now we have all the pieces for e_{S_l} , d_{S_u} and $\text{KL}(Q \| P)$, and can obtain,

$$\begin{aligned}
J(Q) &= C \left[e_{S_l}(G_Q) + \frac{1}{2} d_{S_u}(G_Q) \right] + \frac{1}{m} \text{KL}(Q \| P) \\
&= \frac{1}{2m} \|\mathbf{u} - \mathbf{u}_0\|^2 - \frac{1}{m^2} \sum_i \log Z_i \\
&\leq \frac{1}{2m} \|\mathbf{u} - \mathbf{u}_0\|^2 + \frac{C}{m_l n} \sum_{i=1}^{m_l} \sum_{j=1}^n \Phi(y_i \mathbf{u} \cdot \bar{\phi}_{ij}) \\
&\quad + \frac{C}{m_u n} \sum_{i=1}^{m_u} \sum_{j=1}^n \Phi(\mathbf{u} \cdot \bar{\phi}_{ij}) \Phi(-\mathbf{u} \cdot \bar{\phi}_{ij}) \triangleq J(\mathbf{u})
\end{aligned}$$

ACKNOWLEDGMENT

This work is supported by National Basic Research Program of China (Grant No. 2011CB302203), NSFC (Grant No. 60833009 and 60975012) and Microsoft Research Fellowship.

REFERENCES

- [1] D. Blei, A. Ng, and M. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
- [2] P. Germain, A. Lacasse, F. Laviolette, and M. Marchand. PAC-Bayesian learning of linear classifiers. In *ICML*, pages 353–360, 2009.
- [3] T. Jaakkola and D. Haussler. Exploiting generative models in discriminative classifiers. In *NIPS*, pages 487–493, 1999.
- [4] M. Jordan, Z. Ghahramani, J. Tommi, and S. Lawrence. Introduction to variational methods for graphical models. *Machine Learning*, 37:183–233, 1999.
- [5] J. Langford. Tutorial on practical prediction theory for classification. *Journal of Machine Learning Research*, 6(1):273, 2006.
- [6] D. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- [7] R. Neal and G. Hinton. A view of the EM algorithm that justifies incremental, sparse, and other variants. *Learning in Graphical Models*, 89:355–368, 1998.
- [8] A. Y. Ng and M. I. Jordan. On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes. In *NIPS*, 2002.
- [9] A. Perina, M. Cristani, U. Castellani, V. Murino, and N. Jojic. Free energy score space. In *NIPS*, pages 1428–1436, 2009.
- [10] L. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [11] V. Vapnik. *The nature of statistical learning theory*. Springer Verlag, 2000.
- [12] A. Lacasse, F. Laviolette, and M. Marchand. PAC-Bayes bounds for the risk of the majority vote and the variance of the Gibbs classifier. In *NIPS*, 2006.
- [13] D. McAllester. Some PAC-Bayesian theorems. *Machine Learning*, 37:355–363, 1999.
- [14] X. Li, T. Lee, Y. Liu. Hybrid generative-discriminative classification using posterior divergence. In *CVPR*, 2011.
- [15] T. Griffiths, and M. Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101:5228–5235, 2004.
- [16] A. Vedaldi, V. Gulshan, M. Varma, and A. Zisserman. Multiple Kernels for Object Detection. In *ICCV*, 2009.
- [17] Y. Freund, and R. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Computational Learning Theory*, pages 23–37, 1995.
- [18] M. Gonen, and E. Elpaz. Localized Multiple Kernel Learning. In *ICML*, 2008.
- [19] T. Jaakkola, M. Meila, and T. Jebara. Maximum entropy discrimination. *Technical Report AITR-1668*, MIT, 1999.
- [20] A. Holub, M. Welling, and P. Perona. Hybrid generative-discriminative visual categorization. *International Journal of Computer Vision*, 77(1):239–258, 2008.
- [21] R. Raina, Y. Shen, A. Ng, and A. McCallum. Classification with hybrid generative/discriminative models. In *NIPS*, 2004.
- [22] L. Baum, T. Petrie, G. Soules, and N. Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *The Annals of Mathematical Statistics*, 41(1):164–171, 1970.
- [23] A. McCallum, C. Pal, G. Druck, and X. Wang. Multi-conditional learning: Generative/discriminative training for clustering and classification. In *AAAI*, pages 433–439, 2006.